

SUFFICIENT CONDITION ON NOISE CORRELATIONS FOR SCALABLE QUANTUM COMPUTING

JOHN PRESKILL

*Institute for Quantum Information and Matter, California Institute of Technology
Pasadena, CA 91125, USA*

I study the effectiveness of fault-tolerant quantum computation against correlated Hamiltonian noise, and derive a sufficient condition for scalability. Arbitrarily long quantum computations can be executed reliably provided that noise terms acting collectively on k system qubits are sufficiently weak, and decay sufficiently rapidly with increasing k and with increasing spatial separation of the qubits.

Keywords: Quantum error correction, fault tolerance, accuracy threshold

1 Introduction

Our planet is in the midst of a digital revolution, validating the scalability of classical information processing. Will quantum computers likewise be scalable, eventually performing tasks that surpass what could be done if the world were classical?

The accuracy threshold theorem for quantum computation establishes that scalability is achievable provided that the currently accepted principles of quantum physics hold and that the noise afflicting a quantum computer is neither too strong nor too strongly correlated [1, 2, 3, 4, 5, 6, 7]. For scalability to fail as a matter of principle then, either quantum mechanics must fail for complex highly entangled systems (as 't Hooft [8] has suggested), or else either the Hamiltonian or the quantum state of the world must impose noise correlations that overwhelm fault-tolerant quantum protocols (as Alicki *et al.* [9, 10, 11] and Kalai [12, 13, 14, 15, 16] have suggested).

Because of the profound implications of large-scale quantum computing for computational complexity theory and fundamental physics, skepticism is natural and useful. Debate about the feasibility of fault-tolerant quantum computation can sharpen our understanding of the issues, and raise the stakes as quantum science and technology continue to advance. But skeptics should be pressed for a conception of Nature in which classical computing is feasible yet quantum computing is forbidden.

The theory of fault-tolerant quantum computing closely resembles the corresponding classical theory, but there are also important differences. Perhaps most fundamentally, a classical

computer can perform reliably even if the information being processed leaks to the environment, but a quantum computation will fail unless the processed quantum information remains almost perfectly concealed. Indeed, in classical systems expelling heat to the environment is essential to ensure controllability, while for quantum systems energy dissipation may induce decoherence and hence cause trouble.

A central lesson of fault-tolerant quantum computing is that this is a false dichotomy — energy dissipation is just as crucial for reliable quantum computation as for classical computation [17]. The trick is to expel entropy without exposing the protected coherent quantum information to the environment. In principle this can be achieved using quantum error-correcting codes [18, 19] and carefully designed fault-tolerant protocols [20, 21]. Another important lesson is that small rotation errors in imperfect unitary quantum gates can be digitized and hence corrected like the bit flip errors in a dissipative classical system.

The goal of this paper is to exhibit a class of noise models for which quantum computing is provably scalable. We will assume that qubits can be refreshed on demand, *i.e.*, that it is possible to prepare a standard initial state of n qubits, approximating the product state $|0\rangle^{\otimes n}$, with small, weakly correlated errors. One might imagine that Nature conspires to block the creation of a good approximation to a pure product state, but such a limitation would threaten the scalability of classical computation as well, and so does not seem like a promising way to make a fundamental distinction between classical and quantum computation.

More plausibly, the distinction might arise because noise correlations unavoidably obstruct the creation of profoundly entangled states of many qubits. We will address this issue by studying a class of Hamiltonian models of correlated noise and deriving a sufficient condition for scalability within the context of this class of models. Skeptics are invited to explain why no quantum engineer could ever build a system with noise meeting this criterion.

In our models, the noise correlations arise from an (in principle infinite) series of terms in the Hamiltonian, where terms acting on k system qubits have an operator norm obeying an upper bound that drops sufficiently rapidly with increasing k and with growing spatial separation among the qubits. It would be desirable to relax our scalability criterion in various ways. In particular, scalability can be proven for noise models in which system qubits couple to harmonic oscillator bath variables with unbounded norm, assuming the initial state of the bath meets certain conditions (for example, if the bath starts out in a low-temperature Gibbs state) [22]. But, so far, a proof of scalability for a system in contact with an oscillator bath has been worked out only for Gaussian noise, *i.e.*, for the case where the noise is completely characterized by the two-point correlations function of the bath variables. Extending that argument to a nonlinear oscillator bath with non-Gaussian correlations remains a technically challenging open problem, worthy of further attention.

Gaussian noise models are often regarded as reasonably realistic in physical settings where the system is weakly coupled to many environmental degrees of freedom [23]. But even in the Gaussian case, arguments for scalability hold only under assumptions about the frequency spectrum of bath fluctuations [24, 25, 26, 22, 27], and some skeptics have criticized these assumptions [28, 29, 30]. However these critics have not clearly identified any class of Gaussian noise models which would allow high-fidelity gates in few qubit systems while disallowing large-scale fault-tolerant quantum computing. Hence if their objections carry weight, we may encounter a barrier blocking further systematic improvements in quantum gate fidelity

in the relatively near future. The question we are trying to address here is not whether noise will limit the reliability of small-scale quantum computers but rather whether large scale quantum computing might eventually fail, even though small scale quantum computers continue to improve.

Anyway, without further apologies, we will use a Hamiltonian model in which the noise strength can be characterized using the operator norm. We go beyond previous work [31] by investigating the effects of not just few-body correlations in the noise but also sufficiently weak many-body correlations. The goal is to get a clearer picture how harmful such noise correlations could be. The type of model we study has a notable advantage — we do not need to make any assumption about the initial state of the bath to derive useful results. We formulate the model and state the main result in Sec. 2, then prove it in Sec. 3.

2 Noise model and scalability criterion

The noise model we consider is formulated by specifying a time-dependent Hamiltonian H that governs the joint evolution of the system and the bath, which can be expressed as

$$H = H_S + H_B + H_{SB}; \quad (1)$$

here H_S is the time-dependent Hamiltonian of the system that realizes an ideal quantum circuit, H_B is the Hamiltonian of the bath, and H_{SB} , which describes the coupling of the system to the bath, is the origin of the noise. We place no restrictions on the bath Hamiltonian H_B . Without any loss of generality, we may expand the system-bath Hamiltonian in the form

$$H_{SB} = \sum_{k=1}^{\infty} \sum_{\langle i_1, i_2, \dots, i_k \rangle} H_{i_1, i_2, \dots, i_k}^{(k)} = \sum_{k=1}^{\infty} \frac{1}{k!} \sum_{i_1, i_2, \dots, i_k} H_{i_1, i_2, \dots, i_k}^{(k)}. \quad (2)$$

Here, $H_{i_1, i_2, \dots, i_k}^{(k)}$ acts on the k system qubits labeled by the indices i_1, i_2, \dots, i_k , and also acts arbitrarily on the bath; for each k we sum over all ways of choosing k system qubits. We use $\langle i_1, i_2, \dots, i_k \rangle$ to denote an unordered set of k qubits; by definition, $H_{i_1, i_2, \dots, i_k}^{(k)}$ is invariant under permutations of the k qubits and vanishes if two of the indices coincide. Hence the two expressions for H_{SB} in Eq.(2) are equivalent. We will not need to assume anything about the initial state of the bath, except that the system qubits can be well enough isolated from the bath that we can prepare single-qubit states with reasonable fidelity.

We use the term *location* to speak of an operation in a quantum circuit performed in a single time step; a location may be a single-qubit or multi-qubit gate, a qubit preparation step, a qubit measurement, or the identity operation in the case where a qubit is idle during a time step. We model a noisy preparation as an ideal preparation followed by evolution governed by H , and a noisy measurement as an ideal measurement preceded by evolution governed by H . It is convenient to imagine that all system qubits are prepared at the very beginning of the computation and measured at the very end; in that case the noisy computation can be fully characterized by a unitary evolution operator U acting jointly on the system and the bath, obtained by solving the time-dependent Schrödinger equation for the Hamiltonian Eq.(2).

Let's briefly explain the physical justification for these assumptions. For fault-tolerant computing to work, there must be a mechanism for flushing the entropy introduced by noise; typically, entropy is removed from the computer by error-correction gadgets which use a

supply of fresh ancilla qubits that are discarded after use. For mathematical convenience, we suppose that the initial state of the system includes all of the ancilla qubits that will be needed during the full course of the computation. To model the actual situation, in which ancilla qubits are prepared as needed just before being used, we also suppose that ancilla qubits are perfectly isolated from the bath until “opened” at the onset of the gadget in which they participate. Similarly, we suppose that the measurements of all ancilla qubits are delayed until the very end of the computation, but that these qubits are “closed” (their coupling to the bath is turned off) at the conclusion of the gadget in which they participate. Fault-tolerant gadgets sometimes also include quantum gates that are conditioned on the classical outcomes of earlier measurements. These conditional gates can be included in our framework; operations conditioned on measurements may be replaced by coherent gates, conditioned on the state of a “closed” control qubit that will be measured later (see Sec. VIC of [22]). With these stipulations, our noise model is equivalent to a more realistic one in which ancilla qubits are repeatedly measured, reset, and reused. In this model we take for granted that “pretty good” fresh ancillas can be prepared at any time, or equivalently that qubits can be effectively erased at any time. A similar assumption would be needed to ensure scalability in an analysis of fault-tolerant reversible classical computation [17].

Our goal is to derive from Eq.(2) an expression for the *effective noise strength* ε of the noisy computation, which is defined as follows [6, 31]. We envision performing a formal expansion of U in powers of the perturbation H_{SB} , to all orders. Consider a particular set \mathcal{I}_r of r circuit locations, and let $E(\mathcal{I}_r)$ denote the sum of all terms in the expansion such that every location in \mathcal{I}_r is faulty, *i.e.*, such that at least one of the qubits at that location is struck at least once by a term in H_{SB} during the execution of the gate. We say that the noise has effective noise strength ε if

$$\|E(\mathcal{I}_r)\| \leq \varepsilon^r \quad (3)$$

for any set \mathcal{I}_r . The accuracy threshold theorem for quantum computing shows that scalable quantum computing is possible if ε is less than a positive constant $\varepsilon_0 \approx 10^{-4}$ [6, 22].

Let us define

$$\tilde{\eta}_1^{(k)} = \max_{i_1} \sum_{i_2, i_3, \dots, i_k} \|H_{i_1, i_2, i_3, \dots, i_k}^{(k)}\| t_0, \quad (4)$$

where the maximum is over all system qubits and all times, the sum is over all system qubits, and t_0 is the maximal duration of any location. Then our main result can be stated as follows.

Theorem 1 (Effective noise strength for correlated Hamiltonian noise) *If each quantum gate acts on at most m qubits and if*

$$\tilde{\eta}_1^{(k)} \leq f_k \alpha^k, \quad (5)$$

for all k , then

$$\varepsilon \leq 2m\alpha \exp\left(\sum_{k=1}^{\infty} \frac{g_k}{2k!}\right), \quad (6)$$

where

$$g_k = \sum_{l=0}^{\infty} \frac{(k-1)! f_{k+l} (2\alpha)^l}{(k+l-1)!}. \quad (7)$$

It follows that quantum computing is scalable provided the strength of k -qubit interactions decays sufficiently rapidly with k (so that the sums in Eq.(6) and Eq.(7) converge), and also decays as the spatial separation of the qubits increases (so that the sum defining $\tilde{\eta}_1^{(k)}$ in Eq.(4) converges).

If, for example, $f_k = 1$, then

$$g_k \leq \sum_{l=0}^{\infty} (2\alpha)^l = (1 - 2\alpha)^{-1} \equiv C(\alpha), \quad (8)$$

and hence

$$\varepsilon \leq 2m\alpha \left(e^{(e-1)/2} \right)^{C(\alpha)} \approx 4.72 m\alpha, \quad (9)$$

where the last approximation uses $C(\alpha) \approx 1$ for $\alpha \ll 1$, as is the case if ε is smaller than the threshold value $\varepsilon_0 \approx 10^{-4}$. This observation can be restated as the following corollary:

Corollary 1 *If each quantum gate acts on at most m qubits then the effective noise strength can be expressed as*

$$\varepsilon \leq 2m\alpha \left(e^{(e-1)/2} \right)^{C(\alpha)}, \quad (10)$$

where

$$\alpha = \max_{k \geq 1} \left(\max_{i_1} \sum_{i_2, i_3, \dots, i_k} \|H_{i_1, i_2, i_3, \dots, i_k}^{(k)}\| t_0 \right)^{1/k}, \quad (11)$$

t_0 is the maximal duration of any circuit location, and $C(\alpha) = (1 - 2\alpha)^{-1}$.

If instead

$$f_k \leq k!/k^p \quad (12)$$

where $p \geq 1$, then

$$\begin{aligned} g_k &\leq \frac{k!}{k^p} \left(\sum_{l=0}^{\infty} \frac{(k-1)! k^p f_{k+l} (2\alpha)^l}{(k+l-1)! k!} \right) = \frac{k!}{k^p} \left(\sum_{l=0}^{\infty} \frac{(k-1)! k^p (k+l)! (2\alpha)^l}{(k+l-1)! k! (k+l)^p} \right) \\ &\leq \frac{k!}{k^p} \left(\sum_{l=0}^{\infty} \frac{k^{p-1} (2\alpha)^l}{(k+l)^{p-1}} \right) \leq \frac{k!}{k^p} \left(\frac{1}{1-2\alpha} \right) = \frac{k!}{k^p} C(\alpha). \end{aligned} \quad (13)$$

For $p > 1$ the sum over k in Eq.(6) converges, and hence we obtain a finite expression for ε ; therefore, scalable fault-tolerant quantum computation is achievable for sufficiently small (nonzero) α . We have obtained:

Corollary 2 *If each quantum gate acts on at most m qubits then for any $p > 1$ the effective noise strength can be expressed as*

$$\varepsilon \leq 2m\alpha_p \exp\left(C(\alpha_p) \sum_{k=1}^{\infty} \frac{1}{2k^p}\right), \quad (14)$$

where

$$\alpha_p = \max_{k \geq 1} \left[\frac{k^p}{k!} \left(\max_{i_1} \sum_{i_2, i_3, \dots, i_k} \|H_{i_1, i_2, i_3, \dots, i_k}^{(k)}\| t_0 \right) \right]^{1/k}, \quad (15)$$

t_0 is the maximal duration of any circuit location, and $C(\alpha) = (1 - 2\alpha)^{-1}$.

In particular, if $p = 2$ for example, we find

$$\varepsilon \leq 2m\alpha \left(e^{\pi^2/12}\right)^{C(\alpha)} \approx 4.55 m\alpha, \quad (16)$$

again using $C(\alpha) \approx 1$ to obtain the numerical expression.

3 Proof of Theorem 1

In [31], scalability was proven for the special case in which only the $k = 2$ term in the Hamiltonian is nonzero. To prove Theorem 1 we generalize the ideas used in [31]. We write the system-bath Hamiltonian as

$$H_{SB} = \sum_a H_{SB,a} \quad (17)$$

where a is a shorthand for the indices k , and i_1, i_2, \dots, i_k in Eq.(2). For the sake of conceptual clarity we imagine dividing time into infinitesimal intervals, each of width Δ , and express the time evolution operator for the interval $(t, t + \Delta)$ as

$$U(t + \Delta, t) \approx e^{-i\Delta H} \approx e^{-i\Delta H_S} e^{-i\Delta H_B} \prod_a (I_{SB} - i\Delta H_{SB,a}). \quad (18)$$

(We have omitted terms higher order in $\|H\|\Delta$; strictly speaking, then, to justify Eq.(18) we should regulate the bath Hamiltonian H_B by imposing an upper bound on its norm, then choose Δ small enough so these higher order terms can be safely neglected.) We expand $U(t + \Delta, t)$ as a sum of monomials, where for each value of a either I_{SB} or $-i\Delta H_{SB,a}$ appears; then we obtain the perturbation expansion of the full time evolution operator U over time T by stitching together T/Δ such infinitesimal time evolution operators.

We will refer to the r specified locations in the set \mathcal{I}_r as the “marked locations” and to the remaining locations as the “unmarked locations.” For now, suppose for definiteness that all of the marked locations are single-qubit gates. For any term in the perturbation expansion contributing to $E(\mathcal{I}_r)$ there must be an *earliest* infinitesimal time interval in each of the r marked locations where a term $H_{SB,a}$ acts nontrivially on that qubit. Suppose we fix the infinitesimal time intervals where these earliest “insertions” of H_{SB} occur, and also fix the terms $\{H_{SB,a}\}$ in the system-bath Hamiltonian that act there, but sum over all the terms

in the perturbation expansion acting in other time intervals and on other qubits. Then in between the fixed earliest insertions in the marked locations, the joint evolution of the system and the bath is governed by a modified Hamiltonian

$$H^{(\text{modified})} = H_S + H_B + H_{SB}^{(\text{modified})}, \quad H_{SB}^{(\text{modified})} = \sum_a^{(\text{modified})} H_{SB,a}, \quad (19)$$

where the modified sum excludes any term $H_{SB,a}$ acting nontrivially on any one of the marked locations during any time interval prior to the fixed time of the earliest insertion. The important point is that the time evolution operator in between successive insertions of the perturbation is unitary and hence has unit operator norm. Using the submultiplicative property of the operator norm, then, we conclude that the contribution to $E(\mathcal{I}_r)$ with the earliest insertions at the marked locations fixed has operator norm bounded above by

$$\prod_a^{(\text{earliest})} (\|H_{SB,a}\| \Delta), \quad (20)$$

where the product is over the terms in the system-bath Hamiltonian that act at the earliest insertions. To bound $\|E(\mathcal{I}_r)\|$, we sum over the t_0/Δ time intervals at each location where the earliest insertion may occur, and also sum over all the ways of choosing the term $H_{SB,a}$ that acts at each insertion, obtaining

$$\|E(\mathcal{I}_r)\| \leq \sum_{\{H_{SB,a}\}}^{(\text{insertions})} \prod_a^{(\text{earliest})} (\|H_{SB,a}\| t_0). \quad (21)$$

Summing over the possible intervals for the first insertion turns the factor Δ into the factor t_0 .

Now we have to figure out how to sum over all ways of choosing the terms $\{H_{SB,a}\}$ acting at the earliest insertions inside the r marked circuit locations. Since H_{SB} contains multi-qubit terms, a single term in H_{SB} can simultaneously produce the first insertion at multiple circuit locations occurring in the same time step. Specifically, a single term in $H^{(k)}$ might cause simultaneous faults in j of the r marked locations for any $j \leq k$, if all of these j locations occur in the same time step. We use the term “ j -contraction” to refer to the case where a single term in Eq.(2) produces the first insertion in each of j marked locations. See Fig. 1.

First we find an upper bound on the strength of a one-contraction, the operator norm of the sum of all terms that cause one particular circuit location to be faulty. If the qubit at the marked location carries the label i_1 , the term in the Hamiltonian responsible for the earliest insertion at this location could be any $H_{i_1,j_1,j_2,\dots,j_l}^{(1+l)}$ for $l \geq 0$; here for each $m = 1, 2, \dots, l$ either qubit j_m is unmarked or else qubit j_m is marked but has already been struck by an earlier insertion during the same time step. Hence an upper bound on the strength of the one-contraction is

$$\eta_1 = \sum_{l=0}^{\infty} \eta_1^{(1+l)}, \quad (22)$$




Fig. 1. Contractions occurring during one computational time step, with time flowing vertically downward, where rectangles represent qubits and k dots connected by a horizontal line signify an insertion of a term in $H^{(k)}$, which acts on k qubits. All six of the qubits shown have faults during this time step, arising from a one-contraction, a three-contraction, and a two-contraction. The two-contraction is due to a term in $H^{(6)}$ that actually afflicts all six qubits, but counts as a two-contraction because four of these qubits have already been hit by other contractions earlier in the same time step.

where

$$\eta_1^{(1+l)} = \max_{i_1}^{(\text{in})} \sum_{\langle j_1, j_2, \dots, j_l \rangle}^{(\text{all})} \|H_{i_1, j_1, j_2, \dots, j_l}^{(1+l)}\| t_0 = \max_{i_1}^{(\text{in})} \frac{1}{l!} \sum_{j_1, j_2, \dots, j_l}^{(\text{all})} \|H_{i_1, j_1, j_2, \dots, j_l}^{(1+l)}\| t_0. \quad (23)$$

Here, to obtain an upper bound, we sum each index j_m over all qubits, whether marked or unmarked, and also obtain the factor t_0 by summing over all the infinitesimal time intervals during a single time step. We have also maximized this expression over all possible ways to choose qubit i_1 from among the marked qubits — the superscript “in” in $\max^{(\text{in})}$ indicates that we maximize over only the marked qubits, while the superscript “all” in $\sum^{(\text{all})}$ indicates that we sum over all qubits without any restriction. Of course, our upper point would still be valid were we to relax the restriction and maximize over all qubits, whether marked or not.

Similarly, for $k > 1$, the strength of a k -contraction can be bounded by

$$\eta_k = \sum_{l=0}^{\infty} \eta_k^{(k+l)}, \quad (24)$$

where

$$\eta_k^{(k+l)} = \sum_{\langle i_1, i_2, \dots, i_k \rangle}^{(\text{in})} \sum_{\langle j_1, j_2, \dots, j_l \rangle}^{(\text{all})} \|H_{i_1, i_2, \dots, i_k, j_1, j_2, \dots, j_l}^{(k+l)}\| t_0. \quad (25)$$

Here, for the “in” sum the qubits are restricted to the marked locations and for the “all” sum they may be at either marked or unmarked locations. Note that the upper bound η_1 involves a maximum over marked qubits, while the upper bound η_k for $k > 1$ involves instead a sum over marked qubits; the reason for this distinction is explained in the next paragraph.

By summing all ways of choosing the first insertion in each of r marked locations, we obtain the bound

$$\varepsilon^r \leq \sum_{r_1, r_2, r_3, \dots}^{(r)} \prod_{k=1}^{\infty} \frac{1}{r_k!} (\eta_k)^{r_k}. \quad (26)$$

Here r_k is the number of k -contractions, and the sum $\sum^{(r)}$ is subject to the constraint $\sum_k k r_k = r$. To obtain Eq.(26), we observe that, for $k > 1$,

$$\left(\sum_{\langle i_1, i_2, \dots, i_k \rangle}^{(\text{in})} \sum_{\langle j_1, j_2, \dots, j_l \rangle}^{(\text{all})} \| H_{i_1, i_2, \dots, i_k, j_1, j_2, \dots, j_l}^{(k+l)} \| t_0 \right)^{r_k} \quad (27)$$

contains each way of choosing r_k k -contractions among the r marked locations $r_k!$ times, plus additional nonnegative terms; the factor $1/r_k!$ in Eq.(26) compensates for this overcounting. Furthermore, once all the higher rank contractions have been fixed, the locations where one-contractions occur are completely determined. That is why we defined η_1 by maximizing over i_1 , rather than summing i_1 over all the marked qubits.

In Eq.(22) we have derived a bound on the strength of the noise acting at a single circuit location. We wish to go further and investigate whether the correlations in noise acting collectively on many circuit locations could overcome fault-tolerant protocols, even if the individual gates perform very well. For this purpose, we should relate η_k for $k > 1$ to η_1 . Note that in $\eta_k^{(k+l)}$ for $k > 1$, we can replace the sum over ways to choose k qubits by a sum over all qubits divided by $k!$, and similarly we can replace the sum over the ways to choose l qubits by a sum over all qubits divided by $l!$, obtaining

$$\eta_k^{(k+l)} = \frac{1}{k!} \sum_{i_1, i_2, \dots, i_k}^{(\text{in})} \frac{1}{l!} \sum_{j_1, j_2, \dots, j_l}^{(\text{all})} \| H_{i_1, i_2, \dots, i_k, j_1, j_2, \dots, j_l}^{(k+l)} \| t_0. \quad (28)$$

By summing i_1 over the r marked locations we obtain the bound

$$\eta_k^{(k+l)} \leq r \max_{i_1}^{(\text{in})} \frac{1}{k!} \sum_{i_2, \dots, i_k} \frac{1}{l!} \sum_{j_1, j_2, \dots, j_l}^{(\text{all})} \| H_{i_1, i_2, \dots, i_k, j_1, j_2, \dots, j_l}^{(k+l)} \| t_0; \quad (29)$$

note that we still have a bound if we extend the “in” sum to a sum over all qubits. From Eq.(23) we have

$$\eta_1^{(k+l)} = \max_{i_1}^{(\text{in})} \frac{1}{(k+l-1)!} \sum_{i_2, \dots, i_k}^{(\text{all})} \sum_{j_1, j_2, \dots, j_l}^{(\text{all})} \| H_{i_1, i_2, \dots, i_k, j_1, j_2, \dots, j_l}^{(k+l)} \| t_0. \quad (30)$$

which implies, for $k > 1$,

$$\eta_k^{(k+l)} \leq r \frac{(k+l-1)!}{k! l!} \eta_1^{(k+l)} = \frac{r}{k} \binom{k+l-1}{l} \eta_1^{(k+l)} \leq \frac{r}{k} 2^{k+l-1} \eta_1^{(k+l)}. \quad (31)$$

Hence we find, for $k > 1$,

$$\eta_k \leq \frac{r}{2k} \sum_{l=0}^{\infty} 2^{k+l} \eta_1^{(k+l)}. \quad (32)$$

This is the key inequality that we needed, relating (an upper bound on) the strength of collective noise acting on k circuit locations to a sum over (upper bounds on) contributions to the noise strength for a single location.

Now suppose, as in the hypothesis of Theorem 1, that

$$\eta_1^{(k)} = \frac{1}{(k-1)!} \tilde{\eta}_1^{(k)} \leq \frac{f_k \alpha^k}{(k-1)!}. \quad (33)$$

From Eq.(32) we obtain

$$\eta_k \leq \frac{r g_k}{2k!} (2\alpha)^k, \quad (34)$$

where

$$g_k = f_k + \sum_{l=1}^{\infty} \frac{(k-1)! f_{k+l} (2\alpha)^l}{(k+l-1)!}. \quad (35)$$

Then the bound Eq.(26) becomes

$$\varepsilon^r \leq \sum_{r_1, r_2, r_3, \dots}^{(r)} \prod_{k=1}^{\infty} \frac{1}{r_k!} \left(\frac{r g_k (2\alpha)^k}{2k!} \right)^{r_k} = (2\alpha)^r \sum_{r_1, r_2, r_3, \dots}^{(r)} \prod_{k=1}^{\infty} \frac{1}{r_k!} \left(\frac{r g_k}{2k!} \right)^{r_k}, \quad (36)$$

recalling the constraint on the sum. If we now relax the constraint on the sum, we have

$$\begin{aligned} \varepsilon^r &\leq (2\alpha)^r \prod_{k=1}^{\infty} \sum_{r_k=0}^{\infty} \frac{1}{r_k!} \left(\frac{r g_k}{2k!} \right)^{r_k} = (2\alpha)^r \prod_{k=1}^{\infty} \exp \left(\frac{r g_k}{2k!} \right) \\ &= (2\alpha)^r \left(\exp \left(\sum_{k=1}^{\infty} \frac{g_k}{2k!} \right) \right)^r = \left(2\alpha \exp \left(\sum_{k=1}^{\infty} \frac{g_k}{2k!} \right) \right)^r. \end{aligned} \quad (37)$$

Up until now, we have considered all circuit locations to be single-qubit locations. In the case of an m -qubit gate location, the location is faulty if the system-bath perturbation acts nontrivially on any one of m qubits, which enhances each η_k appearing in Eq.(26) by at most a factor of m^k , and hence increases our upper bound on the effective noise strength by at most a factor of m . This completes the proof of Theorem 1. \square

4 Conclusions and outlook

Theorem 1, combined with results from [6], shows that fault-tolerant quantum computing is scalable in principle for a class of correlated noise models. The key feature of these models is that, while there are terms in the noise Hamiltonian acting collectively on k system qubits for $k \gg 1$ (and simultaneously on the quantum computer's environment, *i.e.*, the “bath”), the operator norm of these terms decays faster than any power of k . The result may apply even if, for each fixed k , the operator norm of the k -qubit noise term decays algebraically rather than exponentially as the qubits are spatially separated, provided the decay is sufficiently rapid as a function of distance for the sum over system qubits in Eq.(4) to converge.

Theorem 1 generalizes results found in [31, 25] for the case $k = 2$. Though the analysis is formulated in terms of interaction-picture perturbation theory, it is rigorous because our estimate of the effective noise strength is derived by bounding perturbation theory summed to all orders.

We also assume that at all times during the computation it is possible to prepare qubits in a standard initial state and to measure qubits in a standard basis with reasonable fidelity, but we make no other assumptions about the state of the bath. Although the bath might be quite “hot,” the flow of entropy from the bath to the computer is impeded by the weak coupling between the system and the bath, allowing a fault-tolerant protocol to maintain a steady state where entropy is removed fast enough to keep the computer “cool.” Entropy is carried away by the ancilla qubits used in error correction gadgets, which are subsequently erased and reused. A mechanism for flushing entropy is required in any scheme for stabilizing a noisy computer, whether quantum or classical [17].

Of course, the condition for scalability derived in Theorem 1 is merely sufficient and not necessary. In particular, scalability may be provable even if the system-bath Hamiltonian has unbounded norm, but in that case further assumptions are needed about the state of the bath at the beginning of the computation. For example, threshold theorems for Gaussian noise were proven in [22], which apply if the bath is a linear system of harmonic oscillators and the initial state of the bath is Gaussian (a thermal state for example). In that case, the criterion for scalability can be stated as a property of the bath’s two-point correlation function, ensuring that spatial and temporal noise correlations decay sufficiently rapidly. In this particular setting, at least, we can address the concern expressed in [9] that over the course of a long computation an oscillator bath could be driven to a highly adversarial state that overwhelms fault-tolerant protocols.

The purpose of this work is to address an issue of principle concerning the scalability of quantum computers subject to correlated noise, not to obtain optimized realistic estimates of the accuracy threshold for quantum computing. Indeed the noise strength ε appearing in our criterion for scalability is in effect an error amplitude per gate rather than an error probability per gate, and must be below about 10^{-4} for known threshold theorems to apply [6, 22]. This criterion is probably much too pessimistic — for the general class of Hamiltonian noise models considered here, we cannot easily rule out substantial enhancement of the logical failure probability due to constructive interference of many coherently combined fault histories, even though this seems quite unlikely in practice. Furthermore, we have not attempted here to assess the effectiveness of methods such as noiseless subsystems [32] or dynamical decoupling [33] which could suppress the noise correlations. Characterizing the residual noise correlations when dynamical decoupling is employed seems difficult for general Hamiltonian noise models, though some preliminary steps were reported in [34]. Also, to derive Theorem 1 we made no assumptions about the bath Hamiltonian H_B , and it may be possible to derive stronger results contingent on physically motivated limitations on the bath dynamics such as locality constraints.

The modest results derived here can hardly be expected to assuage the quantum computing skeptics, but may nevertheless help to clarify the debate. Can we identify fundamental principles of physics that are compatible with large-scale classical computing yet incompatible with large-scale quantum computing? Enlarging the class of noise models for which quantum computing is provably scalable should shed light on this fascinating and important question. If quantum mechanics is valid and if future quantum engineers can devise a controllable many-qubit system with noise meeting the criterion derived in this paper, then reliable large-scale quantum computing will be achievable.

Acknowledgments

I thank Dick Lipton and Ken Regan for allowing me to post a link to a preliminary account of this work on their blog *Gödel's Lost Letter*, and I thank the many readers who posted useful comments on the blog, especially Robert Alicki, Joe Fitzsimons, Aram Harrow, Gil Kalai, and John Sidles. I also thank Peter Brooks and Michael Beverland for discussions. This work was supported in part by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center contract number D11PC20165. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC or the U.S. Government. I also acknowledge support from NSF grant PHY-0803371, DOE grant DE-FG03-92-ER40701, and NSA/ARO grant W911NF-09-1-0442. The Institute for Quantum Information and Matter (IQIM) is an NSF Physics Frontiers Center with support from the Gordon and Betty Moore Foundation.

References

1. D. Aharonov and M. Ben-Or, Fault-tolerant quantum computation with constant error, Proc. 29th Ann. ACM Symp. on Theory of Computing, p. 176 (New York, ACM, 1998), arXiv:quant-ph/9611025; D. Aharonov and M. Ben-Or, Fault-tolerant quantum computation with constant error rate, arXiv:quant-ph/9906129 (1999).
2. A. Yu. Kitaev, Quantum computations: algorithms and error correction, Russian Math. Surveys 52, 1191-1249 (1997).
3. E Knill, R. Laflamme, W. H. Zurek, Resilient quantum computation: error models and thresholds, Proc. Roy. Soc. London, Ser. A 454, 365 (1998), arXiv:quant-ph/9702058.
4. J. Preskill, Reliable quantum computers, Proc. Roy. Soc. Lond. A 454, 385-410 (1998), arXiv:quant-ph/9705031.
5. D. Gottesman, Stabilizer codes and quantum error correction, Caltech Ph.D. thesis (1997), arXiv:quant-ph/9705052.
6. P. Aliferis, D. Gottesman, and J. Preskill, Quantum accuracy threshold for concatenated distance-3 codes, Quant. Inf. Comp. 6, 97-165 (2006), arXiv:quant-ph/0504218.
7. B. W. Reichardt, Threshold for the distance three Steane quantum code, arXiv:quant-ph/0509203 (2005).
8. G 't Hooft, Quantum gravity as a dissipative deterministic system, Classical and Quantum Gravity 16, 3263-3279 (1999), arXiv:gr-qc/9903084
9. R. Alicki, M. Horodecki, P. Horodecki, and R. Horodecki, Dynamical description of quantum computing: generic nonlocality of quantum noise, Phys. Rev. A 65, 062101 (2002), arXiv:quant-ph/0105115.
10. R. Alicki, Quantum error correction fails for Hamiltonian models, Fluctuation and Noise Letters 6, C23-C28 (2006), arXiv:quant-ph/0411008.
11. R. Alicki, Quantum memory as a perpetuum mobile of the second kind, arXiv:0901.0811 (2009).
12. G. Kalai, Thoughts on noise and quantum computation, arXiv:0508095 (2005).
13. G. Kalai, How quantum computers can fail, arXiv:0607021 (2006).
14. G. Kalai, Detrimental decoherence, arXiv:0806.2443 (2008).
15. G. Kalai, Quantum computers: noise propagation and adversarial noise models, arXiv:0904.3265 (2009).
16. G. Kalai, How quantum computers fail: quantum codes, correlations in physical systems, and noise accumulation, arXiv:1106.0485 (2011).

17. D. Aharonov, M. Ben-Or, R. Impagliazzo, and N. Nisan, Limitations of noisy reversible computation, arXiv:quant-ph/9611028 (1996).
18. P. W. Shor, Scheme for reducing decoherence in quantum computer memory, Phys. Rev. A 52, 2493 (1995).
19. A. Steane, Error-correcting codes in quantum theory, Phys. Rev. Lett. 77, 793 (1996).
20. P. W. Shor, Fault-tolerant quantum computation, in Proceedings, 37th Annual Symposium on Foundations of Computer Science, pp. 56-65 (Los Alamitos, CA, IEEE Press, 1996), arXiv:quant-ph/9605011.
21. D. Gottesman, An introduction to quantum error correction and fault-tolerant quantum computation, arXiv:0904.2557 (2009).
22. H.-K. Ng and J. Preskill, Fault-tolerant quantum computation versus Gaussian noise, Phys. Rev. A 79, 032318 (2009), arXiv:0810.4953.
23. A. O. Caldeira and A. J. Leggett, Quantum tunneling in a dissipative system, Ann. Phys. 149, 374 (1983).
24. B. M. Terhal and G. Burkard, Fault-tolerant quantum computation for local non-Markovian noise, Phys. Rev. A 71, 012336 (2005), arXiv:quant-ph/0402104.
25. E. Novais, E. R. Mucciolo, H. U. Baranger, Resilient quantum computation in correlated environments: A quantum phase transition perspective, Phys. Rev. Lett. 98, 040501 (2007), arXiv:quant-ph/0607155.
26. E. Novais, E. R. Mucciolo, H. U. Baranger, Hamiltonian formulation of quantum error correction and correlated noise: The effects of syndrome extraction in the long time limit, Phys. Rev. A 78, 012314 (2008), arXiv:0710.1624.
27. E. Novais, E. R. Mucciolo, and H. U. Baranger, Bound on quantum computation time: Quantum error correction in a critical environment, Phys. Rev. A 82, 020303(R) (2010), arXiv:1004.3247.
28. M. I. Dyakonov, Is fault-tolerant quantum computation really possible?, arXiv:0610117 (2006).
29. R. Alicki, Comment on 'Resilient Quantum Computation in Correlated Environments: A Quantum Phase Transition Perspective' and 'Fault-tolerant Quantum Computation with Long-range Correlated Noise,' arXiv:quant-ph/0702050 (2007).
30. A. P. Hines and P. C. E. Stamp, Decoherence in quantum walks and quantum computers, Canadian Journal of Physics 86, 541-548 (2008), arXiv:0711.1555.
31. D. Aharonov, A. Kitaev, and J. Preskill, Fault-tolerant quantum computation with long-range correlated noise, Phys. Rev. Lett. 96, 050504 (2006), arXiv:quant-ph/0510231.
32. E. Knill, R. Laflamme, and L. Viola, Theory of quantum error correction for general noise, Phys. Rev. Lett. 84, 2525-2528 (2000), arXiv:quant-ph/9908066.
33. L. Viola, E. Knill, and S. Lloyd, Dynamical decoupling of open quantum systems, Phys. Rev. Lett. 82, 24172421 (1999), arXiv:quant-ph/9809071.
34. H.-K. Ng, D. A. Lidar, and J. Preskill, Combining dynamical decoupling and fault-tolerant quantum computation, Phys. Rev. A 84, 012305 (2011), arXiv:0911.3202.